

Robotics Research Technical Report

generatorium omnis laboris ex machina

Analysis of the Motion-Planning
Problem for a Simple
Two-Link Planar Arm

by

Boris Aronov
C. Ó'Dúnlaing

Technical Report No. 314
Robotics Report No. 118
August, 1987


New York University

Department of Mathematical Sciences

Computer Science Division

251 Mercer Street New York, N.Y. 10012

NYU COMPSCI TR-314
Aronov, Boris
Analysis of the motion-
planning problem for a
simple two-link planar arm
c.1



-

Analysis of the Motion-Planning
Problem for a Simple
Two-Link Planar Arm

by

Boris Aronov
C. Ó'Dúnlaing

Technical Report No. 314
Robotics Report No. 118
August, 1987

New York University
Dept. of Computer Science
Courant Institute of Mathematical Sciences
251 Mercer Street
New York, New York 10012

Work on this paper has been supported by Office of Naval Research Grant N00014-82-K-0381, National Science Foundation CER Grant DCR-83-20085, and by grants from the Digital Equipment Corporation and the IBM Corporation.

Analysis of the motion-planning problem for a simple two-link planar arm

Boris Aronov

Colm Ó'Dúnlaing¹

Courant Institute of Mathematical Sciences
251 Mercer Street
New York, New York 10012

ABSTRACT

A two-link planar robot arm is considered inside a planar polygonal workspace. The configuration space of the arm is shown to have $O(n^3)$ components. Modeled on the Leven-Sharir algorithm, a *connectivity graph* CG is constructed, of size $O(n^3)$, in time $O(n^3 \log(n))$. This graph can be used to solve any instance of the motion-planning problem for the arm.

August 9, 1987

¹This work was supported in part by the NSF under grant DCR-84-01898. Ó'Dúnlaing's current address: School of Mathematics, Trinity College, Dublin 2, Irish Republic. Some of this work was reported at the SLAM conference on Applied Geometry, July 1987, Albany, New York.

1. Introduction.

This paper addresses the motion-planning problem for a 2-link robot 'arm' in a planar workspace with polygonal obstacles. Our intention is to adapt the known motion-planning algorithm of Leven and Sharir [LS] to this problem. The idea is quite simple: the Leven-Sharir algorithm is for moving an oriented line-segment of fixed length (humorously called a 'ladder') around a planar 'room' amid polygonal obstacles. We observe that a simple two-link robot arm can be viewed, fairly realistically, as having an anchored *shoulder*, to which is attached an extensible *upper arm*, which is in turn attached to a fixed-length *forearm*. The joint connecting these two links we call the *elbow*. For convenience we shall assume that the upper arm can vary in length from zero to infinity -- or at least be sufficiently long to reach any part of the room (ignoring obstacles). We shall assume that both joints can rotate through 360° ; in a concluding section we shall discuss briefly how some of these assumptions can be relaxed.

1.1. The particular robot model used in our analysis.

We will consider a two-link robot anchored at a fixed origin O . Its telescopic upper arm is modeled by a straight-line segment OP of variable length which is free to rotate around the *shoulder* O . Its forearm PQ is of fixed length d and can rotate freely about the elbow P . It is of course assumed that the origin itself is located inside the room but outside any obstacle.

This problem formulation is analogous to that of [LS] in that the forearm is similar to the ladder. However, an additional constraint is introduced as the upper arm OP is to remain clear of the obstacles. A simplifying assumption is made that the extent of the upper arm is not limited, which would in practice mean that it is great enough to allow the elbow to reach the outside room walls. The Leven-Sharir algorithm is modified to deal with what we call 'virtual walls,' representing constraints imposed by the requirement that the upper arm is not to intersect obstacles. Before describing the algorithm we analyze the structure of the set of allowable motions of the robot arm. Our analysis largely parallels that in [LS] with some

modifications to accommodate 'virtual walls.'

2. Analysis.

Define the *forearm* PQ to be an oriented segment of length d^2 , with *elbow* P . A *placement* of the arm is a pair $Z = (X, \theta)$, where X is the position of P and θ is the orientation of PQ . The *free room* V is the open polygonally bounded region given by room interior less the obstacles. A *real wall* is a boundary segment of V . The *workspace* F is the largest open star-shaped region with center at O contained in V . F is clearly polygonally bounded and is cut out of V by segments each originating at an obstacle corner and directed away from O . We refer to these segments as *virtual walls*; from now on 'wall' will mean a real or virtual wall. Clearly, P must stay within F since the line-segment OP (the upper arm) must stay within V .

Definition. The *set of free placements* FP (respectively, *semi-free placements* SFP) is the set of all placements of the robot arm such that both OP and PQ are completely contained in V (respectively, in the closure of V).

Note. We shall assume that SFP is the closure of FP , which means that SFP cannot include placements in which the arm touches obstacles without penetrating them but any small movement of the arm makes it penetrate obstacles. This simplifying assumption can be made if we require the obstacles to be 'in general position' in the sense that, if such placements exist in SFP , then they can be removed by an arbitrarily small perturbation of the positions of the obstacles.

2.1. Purely translational motion.

For the purposes of this section, let θ be a fixed orientation, which we shall visualize as being directed vertically upwards.

² We will loosely follow the notation of [LS] with *forearm* replacing *ladder* as appropriate.

Definition. Given a placement $Z = (X, \theta) \in FP$, let the *range at Z* (or *range at X with respect to orientation θ*) $I(Z)$, be the maximal open vertical segment through X traced out by P as the forearm moves in FP along a vertical line through X . In other words, it is the set of all elbow positions in placements reachable from Z by a vertical motion of the forearm. The *clearance segment $J(Z)$* at Z is the set of points swept out by the forearm as P varies over the range at Z , i.e., $J(Z)$ is $I(Z)$ extended upwards by a segment of length d .

Definition. The *upper stop* of $Z = (X, \theta)$ (or the upper stop of X with respect to an orientation θ) is the real wall which Q touches or the virtual wall which P touches when P is at the upper end of $I(Z)$; similarly the *lower stop* is the real or virtual wall which P touches when it is at the lower end of $I(Z)$ (see Figure 1). If the upper end of $J(Z)$ is at a corner of the free room V , or its lower end is at a corner of the workspace F , the stop consists of the pair of walls defining that corner. In the case where at the upper end of $I(Z)$ both Q touches a real wall and P touches a virtual one, the pair of these walls is considered to be the upper stop.

Since all lines containing virtual walls intersect at the origin, at most one virtual wall can delimit the upper (respectively, lower) end of $I(Z)$, so at most one virtual wall can participate in the upper (respectively, lower) stop. Henceforth a corner of the room V will be called a *real corner*, and any other corner of F a *virtual corner* (such a corner is at the intersection of a real wall and a virtual wall).




Figure 1.

Definition. The *label* $L(Z)$ of $Z = (X, \theta) \in FP$ is defined to be

$$L(Z) = (\text{lower stop}(Z), \text{upper stop}(Z)).$$

Note that all points on $I(X, \theta)$ share the same label (with respect to θ). A placement $Z = (X, \theta) \in FP$ is *generic* if *both* the upper and the lower stops of Z are defined by *exactly* one wall.

Lemma 2.1. Given a generic placement $Z = (X, \theta)$, there is an open neighborhood $U \subset \mathbb{R}^2$ of X on which L is constant, i.e., $(\forall X' \in U) L(X', \theta) = L(Z)$. Note that implicitly $Z \in FP$, so $|J(Z)| > d$.

The simple proof is omitted: the result holds because a generic placement is one which does not change its label with a sufficiently small horizontal motion, while a small vertical motion does not change the label of any free placement by the above observation.

Definition. Let $Z = (X, \theta) \in FP$ be generic. The (*maximal noncritical*) *region* of Z , denoted $R(Z)$, is the connected component of the set

$$\{ X' \mid (X', \theta) \in FP \ \& \ L(X', \theta) = L(X, \theta) \}$$

containing X .

Clearly, a cross section of such a region by a line with orientation θ through point X' is just $I(X', \theta)$, while the left and right limits of the region are determined by non-generic placements or by reaching the boundary of FP . Moreover, a maximal noncritical region is triangular or trapezoidal in shape and in particular is (the interior of) a convex polygon (for an example, see Figure 2). The placements that define the left and right boundaries of a maximal noncritical region are either non-generic or non-free, with the following cases possible for such placements Z :

- (a) the upper/lower end of $J(Z)$ is at a real corner or its lower end is at a virtual corner,
- (b) the length of the clearance segment $J(Z)$ is exactly d ,
- (c) a corner comes to touch the interior of $J(Z)$, at which point $J(Z)$ in general splits into two pieces, or




Figure 2.

- (d) a virtual wall that intersects $J(Z)$ and defines the upper endpoint of the range $I(Z)$ at Z comes to be at a distance of exactly d from a real wall directly above it.

Case (d) differs from (a) in that there need not be actual intersection of the virtual and the real walls that define the upper stop for Z , but the two nevertheless determine a stop. In cases (a) and (d) the placement Z is non-generic, in case (b) it is in the boundary of FP , and in case (c) both are possible. Figure 3 illustrates the possible placements that give rise to a region boundary condition. The vertical line segment shows $J(Z)$ in cases (a) (excepting the last subcase) and (d). In (b) and (c), it indicates the limit of $J(Z)$ as the forearm approaches the region boundary from inside the region R , which is on the right. Other solid lines show real walls, while dashed lines indicate virtual walls and dot-dashed lines indicate lines traced by P as Q follows a real wall. Interiors of obstacles and of the region R are shaded.

Definition. The *left (right)* stop of a region R is defined to be the set of walls that cause one of the conditions (a)-(d) to become true at the left (respectively right) boundary of the region.

One can classify the cases of Figure 3 as follows. Let l be the line through X with orientation θ , where X is a point on the left (right) boundary of the region. Then (a), (c), and the last case of (b) involve l passing through a real or a virtual corner, and the two walls meeting at the corner define the (left) stop, while (b) (excepting the last case) and (d)




Figure 3(a). One end determined by a corner.




Figure 3(b). $|J(Z)| = d$, in the last subcase – virtual upper and real lower stop. See also last subcase of 3(a).




Figure 3(c). A corner meeting the interior of $J(Z)$.




Figure 3(d). Interaction of an upper real and an upper virtual stop.

involve l intersecting two walls at points exactly d apart, and these two walls define the left stop in these cases.

Definition. The *label* of a region R is the quadruple consisting of its upper, lower, left, and right stops. The following lemma is an easy consequence of the case-analysis illustrated in Figure 3. (In part (b) of the lemma, perturbation arguments can be used to show that a

free placement not in any (open) maximal noncritical region is in the closure of some such region. This can be false if the obstacles are not in general position.)

Lemma 2.2. For any fixed orientation θ , (a) each (maximal noncritical) region R is characterized uniquely by its label, and consequently there are only finitely many such regions; (b) every (semifree) placement is in the closure of some maximal noncritical region.

□

For any fixed orientation θ we can estimate the number of maximal noncritical regions by bounding the number of times conditions (a)-(d) can arise:

- (a) This can arise no more than $O(n)$ times—a given corner can be an endpoint of at most 2 segments with orientation θ . There are n real corners and $O(n)$ virtual ones.
- (b) The last subcase of (b) can occur as many times as there are virtual corners, or $O(n)$ times. As for the other subcases, observe that for each pair of walls, for a given orientation there is at most one semifree placement where the forearm touches both of the walls. Clearly, such boundary conditions will occur at most $O(n^2)$ times. Figure 4 shows that $\Omega(n^2)$ occurrences are possible, so this bound is tight.




Figure 4.

- (c) Note that this can only happen at a convex corner, while the intersection of a real wall with a virtual wall always creates a reflex corner. Therefore (c) occurs only at convex real corners. However, the sets $J(Z)$ for different regions can overlap. So the question is: how many $J(Z)$'s can there be that pass through the same corner? Answer: $O(n)$, for each has to have an intersection of l with a wall as its lower end and no two can share a lower end. Since a line of fixed orientation through a fixed corner can have at most $O(n)$ intersections with walls, (c) can occur only $O(n)$ times per corner, for a total of $O(n^2)$ times. Figure 5 shows that this bound is tight.
- (d) Similar to (b). For each pair of walls, where one is virtual and one is real, there is at most one semifree placement (at orientation θ) where P is on the virtual wall and Q is on the the real one. Hence (d) occurs $O(n^2)$ times and again this bound is tight (Figure 4).

Definition. The *free boundary* of a region $R = R(X, \theta)$ is




Figure 5.

$$b(R) = \{X' \in \partial(R) \mid (X', \theta) \in FP\},$$

where $\partial(R)$ is the topological boundary of R as a planar set.

Lemma 2.3. $b(R)$ consists of finitely many open vertical segments.

Proof. It is clear that $b(R)$ is a finite collection of relatively open subsets of a trapezoid boundary, for it is the intersection of a trapezoid boundary $\partial(R)$ and an open polygonally bounded region $\{X' \mid (X', \theta) \in FP\}$. Moreover from the definition of upper and lower stops it is clear that no point interior to the non-vertical part of the boundary of R can be in $b(R)$. Since the intersection of the boundary of R with the complement of FP is closed, the corners of R are not in FP either. \square

Lemma 2.4. Let R and R' be distinct (maximal noncritical) regions. If $b(R)$ and $b(R')$ have a point X in common, then the connected components of $b(R)$ and $b(R')$ containing X are identical and equal to $I = I(X, \theta)$.

Proof. Let L be the component of $b(R)$ containing X . By definition, $L \subseteq I$. Suppose that $L \neq I$; I must contain at least one endpoint Y of L . Since the $\partial(R)$ is closed, it contains Y , and therefore Y is in $\partial(R) \cap FP$; it follows that Y is in L , a contradiction: therefore $I = L$. Similarly, I is the component of $b(R')$ containing X . \square

Lemma 2.5. Let R and R' be distinct maximal noncritical regions. Let $X \in R$ and $X' \in R'$. Then $b(R) \cap b(R') \neq \emptyset$ if and only if there exists a continuous orientation-preserving motion of the forearm from (X, θ) to (X', θ) that does not enter a third region.

Proof. It is a straightforward exercise to verify that if R and R' are distinct and their free boundaries intersect then either the left free boundary of R' meets the right free boundary of R or vice-versa: suppose that Y is a point in this intersection. Clearly there is a translational motion from (X, θ) to (Y, θ) and thence to (X', θ) , where all placements are in the union of the closures of R and R' , and since these closures plainly cannot intersect a third region (regions are by definition open) the 'if part' is proved.

To prove the converse, we need to show that if such a motion exists then it must be confined to the union of the closures of R and R' . The set of points X in FP (or rather, points X such that vertical placements with P located at X are in FP) not interior to a third region is comprised of the union of the closures of R and R' together with the union of all free boundary components disjoint from this closure. However, by Lemmas 2.3 and 2.4 the latter union consists of open vertical line-segments whose closures are all disjoint from the closures (in FP) of R and R' , and clearly they cannot be reached from R or R' by an orientation-preserving path in FP not going through any other region. \square

Definition. The *connectivity graph* for a fixed orientation θ , CG_θ , is an undirected graph with maximal noncritical regions as vertices and edges defined as follows: $\{R_1, R_2\}$ is an edge of CG_θ if and only if $b(R_1) \cap b(R_2) \neq \emptyset$.

Lemma 2.6. A generic placement (X', θ) is reachable from a generic placement (X, θ) by a continuous orientation-preserving (i.e., translational) motion if and only if $R' = R(X', \theta)$ is reachable from $R = R(X, \theta)$ in CG_θ .

Proof. (1) \Leftarrow An immediate consequence of Lemma 2.5

(2) \Rightarrow Let π be a suitable path and suppose without loss of generality that $R \neq R'$. Assume that R and R' are in different components of the connectivity graph: we shall argue a contradiction. Suppose C is the component of R , i.e., the set of all nodes in CG_θ reachable from R in the graph. Let (Y, θ) be the first point on the translational path π such that Y is in the closure of some region R'' not in C : such a point exists since the set of path points in R'' is closed. However, Y is also on the free boundary of some region U belonging to C , and R'' is adjacent to U , (by the above lemma), so U is in C , a contradiction. \square

2.2. Unrestricted motion of the robot arm.

In this section we shall consider general (not just translational) motions of the arm, and investigate the dependence of CG_θ on θ .

Definition. An orientation θ is *critical* (or a *criticality*) if there is a region $R(X, \theta)$ the label of whose left (or right) stop is ambiguous in the sense that *two* independent events of the type described in Figure 3 occur simultaneously at the left (or right) stop of $R(X, \theta)$. We shall label θ by the two events. $R(X, \theta)$ is thereafter referred to as a *witness region* (or simply a *witness*) of the criticality θ .

We refer to the number of witnesses as the *complexity* of the critical orientation θ . Usually a criticality has more than one witness, since it arises from the interaction of several regions; as we shall see later, the complexity of a given critical orientation θ can be unbounded (i.e., $\Omega(n)$) in number, and one cannot rule this out by assumptions of general position, as one could in the 'ladder' problem. The criticalities will be classified below. We shall see that they arise from coincidences such as, for instance, the orientation being parallel to a line joining two corners; under assumptions of general position exactly two corners should determine such a criticality. Criticalities of high complexity can arise because there may be several overlapping segments of the form $J(X, \theta)$. If this happens, all these segments are contained in the same open line-segment in the room V (ignoring the visibility constraint), and we conclude that they can only be split up by virtue of virtual walls crossing this open line-segment. This can give rise to $\Omega(n)$ witnesses (worst case). As we will see later, there are $O(n^2)$ critical orientations of complexity $O(n)$ and $O(n^3)$ of complexity $O(1)$.

Note. There are some necessary conditions we can establish for a region R to be a witness to the criticality of an orientation θ . These are somewhat more elaborate than those for the 'ladder problem' given in [LS]. Supposing without loss of generality that θ arises from an ambiguity in the left stop of R , let e be the open left boundary segment of R , let W be obtained by extending R upwards by a vertical distance d (adding a parallelogram to the top of R), and let e' be the open left boundary segment of W : it is just e extended upwards by a segment of length d . Clearly W is in V : its upper boundary is defined by the upper stop of R . Both e and e' may meet the boundary of the workspace F . By definition, R cannot intersect

any real or virtual walls, and W cannot intersect any real walls: but the parallelogram by which it extends R may intersect some virtual walls. Therefore e cannot penetrate any real or virtual walls and e' cannot penetrate any real walls and can cross virtual walls only in its top-most segment of length d .

Lemma 2.7. Suppose R is a maximal noncritical region at an orientation θ , where R is not a witness region (whether or not θ is critical). Then there exists an open interval containing θ such that for each θ' in this interval there is a unique region $R_{\theta'}$ whose label coincides with that of R . Moreover, $R_{\theta'}$ varies continuously (in a sense which can easily be made precise³) with θ' as θ' varies over this interval. There is a maximal interval (θ_1, θ_2) with this property, and both orientations θ_i are critical and distinct. \square

This can be proven by a straightforward case analysis which the interested reader is invited to complete.

Definition. Let R, θ be as above. Define a *cell* to be

$$C = \{ (X', \theta') \mid X' \in R_{\theta} \text{ and } \theta' \in (\theta_1, \theta_2) \},$$

where (θ_1, θ_2) is the maximal open interval around θ as in Lemma 2.7. We shall call θ_1 and θ_2 respectively the *lower and upper θ -bounds* of C . The open interval (θ_1, θ_2) will also be called the *range of C* . A cell defined in terms of $R(Z)$ will be denoted $C(Z)$ (and called 'the cell containing Z ').

Lemma 2.8. Every such cell C is open.

Proof. Consider a placement $Z = (X, \theta)$ in C , where $R(Z)$ is not a witness to the criticality of θ , and let ϵ be the distance from X to the boundary of $R(Z)$. By Lemma 2.7 the boundary segments of $R_{\theta'}$ vary continuously with θ' . Thus there exists $\delta > 0$ such that for all θ' , if $|\theta' - \theta| < \delta$, then the ball $B_{\epsilon/2}(X)$ of radius $\epsilon/2$ centered at X contains no points of the boundary of $R_{\theta'}$. In other words, $B_{\epsilon/2}(X)$ is entirely in $R_{\theta'}$. Therefore,

³For instance, the boundary of the regions varies continuously with respect to the Hausdorff metric.

$$B_{\epsilon/2}(X) \times (\theta - \delta, \theta + \delta) \subset C(Z). \quad \square$$

Lemma 2.9. Given two placements Z and Z' in the same cell C , there exists a motion connecting them in C , and such motions can be defined in a canonical way. In particular, every cell C is path-connected.

Proof. (Sketch). By definition, the cross-section R_θ of C at any orientation θ within its range is path-connected, so it is enough to identify a canonical motion between canonical placements Z_θ in the regions R_θ (technically, in $R_\theta \times \{\theta\}$). Following [LS] we choose Z_θ as the 'centroid' of R_θ — that placement where P is the midpoint of the line connecting the midpoints of the upper and lower boundary segments of R_θ . \square

Definition. The *free boundary* of a cell C is the (topological) boundary of the cell in FP , i.e., $b(C) = \partial(C) \cap FP$. Let $C = \{ (X, \theta) \mid X \in R_\theta, \theta \in (\theta_1, \theta_2) \}$ be a cell. Define

$$\Phi_1(C) = \lim_{\theta \rightarrow \theta_1^-} \text{closure}(R_\theta),$$

$$\Phi_2(C) = \lim_{\theta \rightarrow \theta_2^-} \text{closure}(R_\theta),$$

Also let

$$b_\partial(C) = \{ (X, \theta) \mid X \in b(R_\theta), \theta \in (\theta_1, \theta_2) \},$$

$$b_1(C) = \text{interior}(\Phi_1(C)) \times \{\theta_1\},$$

$$b_2(C) = \text{interior}(\Phi_2(C)) \times \{\theta_2\}.$$

Lemma 2.10. $\Phi_1(C)$ ($\Phi_2(C)$) is a point, a line segment, a trapezoid, or a triangle. Moreover, in the latter two cases, $\text{interior}(\Phi_1(C))$ (respectively, $\text{interior}(\Phi_2(C))$) is either a maximum noncritical region with respect to θ_1 (respectively θ_2) or two such regions with identical upper and lower stops together with the free boundary segment between them.

Proof. (Sketch). Without loss of generality we consider only $\Phi_1(C)$. The first part of the lemma is clear, as $\Phi_1(C)$ is a limit of a continuous family of trapezoids or triangles and therefore a possibly degenerate trapezoid. Let R denote the interior of $\Phi_1(C)$. If R is nonempty then we must argue that it is either a maximal noncritical region or the union of

two such regions together with their separating free boundary. Let w_i ($i=1, 2$) be the upper and lower stops, respectively, associated with C . By continuity arguments it is clear that both these walls (whether real or virtual) define upper and lower stops for all placements in R , and $R \subseteq FP$ (i.e., $R \times \{\theta_1\} \subseteq FP$). An inspection of Figure 3 will show that all the conditions defining left (and right) stop for the translational regions R_θ will also hold in the limit as $\theta \rightarrow \theta_1$. Thus, if all placements (X, θ_1) with X in R have w_1, w_2 only as upper and lower stops, R is clearly a maximal noncritical region. Otherwise, let X be a placement in R which has some other wall in its upper (lower) stop besides w_1 (w_2). A simple case-analysis based on Figure 3 will show that the only possibility is as illustrated in Figure 6: the line at orientation θ_1 through x passes through the upper stop — which must be a virtual wall — and meets an obstacle point at distance d beyond this point of intersection. This introduces a line of orientation θ_1 which splits R into two. By continuity arguments, if we assume that no two walls are parallel, the given obstacle point must be a real wall-corner. By assumptions of general position such a splitting line cannot occur twice for the same orientation θ_1 , and it clearly introduces left and right stops for two regions whose union forms R (see the note below for a clarification). This concludes the argument. \square




Figure 6.

Note. For each real corner p , there are $O(n)$ orientations θ_1 such that the line at orientation θ_1 through p meets a virtual wall at distance d . By an arbitrarily small perturbation of the obstacle corners if necessary, we can ensure that for any two different corners these sets of orientations do not intersect. Hence it is legitimate to assume that the cell R is split by at most one line as discussed in the lemma. Moreover, when the region R considered in the above lemma is split as described, we can assume (another general position assumption) that the upper virtual wall w_1 is not perpendicular to θ_1 , i.e., that the given corner which causes the splitting is not at distance exactly d from the virtual wall. With this assumption the two maximal noncritical regions whose union (together with the common free boundary) forms $\Phi_1(C)$ are associated with two non-empty cells whose upper θ -bound is θ_1 .

Corollary 2.11. For $i=1$ or 2 , if $\Phi_i(C)$ has nonempty interior, then it is either a witness region for θ_i or two witness regions separated by the line on which the criticality of θ_i is manifested. All placements in its interior (or more precisely where the elbow is in its interior and the orientation is θ_i) are free placements. \square

Corollary 2.12. For $i=0,1,2$, $b_i(C) \subset FP$.

Proof. True by definition for $i=0$ and follows from Corollary 2.11 for the other two cases. \square

Lemma 2.13. Every placement in $b_0(C) \cup b_1(C) \cup b_2(C)$ is in $b(C)$, and the rest of $b(C)$ is the set of free placements of the form (X, θ_i) with $X \in \partial(\Phi_i(C))$ for $i=1,2$. This remaining set of placements is just $b(\Phi_i(C))$ if $\Phi_i(C)$ is a maximal noncritical region; in any case it is a finite union of open vertical line-segments.

Proof. Consider a placement (X, θ) with $\theta \in [\theta_1, \theta_2]$ and $X \notin \text{closure}(R_0)$ if $\theta \neq \theta_1, \theta_2$ and $X \notin \Phi_i(C)$ if $\theta = \theta_i$. Analogous to the proof of Lemma 2.8, if $\epsilon > 0$ is the distance of X from the boundary of R_0 (or of $\Phi_i(C)$, as appropriate), there exists some $\delta > 0$ such that $B_{\epsilon/2}(X) \times (\theta - \delta, \theta + \delta) \cap C = \emptyset$; thus $(X, \theta) \notin \text{closure}(C)$. Therefore, $\text{closure}(C)$ is confined to $\text{closure}(\Phi_i(C)) \times \{\theta_i\}$, for $i=1,2$, and $\text{closure}(R_0) \times \{\theta\}$, for $\theta \in (\theta_1, \theta_2)$. It is evident,

therefore, that if (X, θ) is a placement in $b(C)$ where θ is in the (open angular) range of C then $X \in b(R_\theta)$, and if it is in $b(C)$ with $\theta = \theta_1$, say, then clearly either $(X, \theta) \in b_1(C)$ or $X \in (FP \cap \partial(\Phi_1))$. This essentially concludes the proof: it is straightforward to verify that the description of $b(C) \sim (b_0(C) \cup b_1(C) \cup b_2(C))$ is correct. \square

Corollary 2.14. The placements in $b(C) \sim (b_0(C) \cup b_1(C) \cup b_2(C))$ constitute a 1-dimensional subvariety of the three-dimensional manifold FP , and their removal cannot disconnect FP nor any open subset of FP .

We can therefore ignore such placements in the motion-planning problem. The proof of the above corollary follows from general topological considerations (see [SS1],[SS2]).

Definition. Two cells C_1 and C_2 are *adjacent* if and only if

$$(b_0(C_1) \cup b_1(C_1) \cup b_2(C_1)) \cap (b_0(C_2) \cup b_1(C_2) \cup b_2(C_2)) \neq \emptyset.$$

Some of the nine pairwise intersections are always empty: we proceed to establish which of them can be nonempty.

Lemma 2.15. For any pair of distinct cells C_1 and C_2 , (a) $b_0(C_1) \cap b_1(C_2) = \emptyset$; (b) $b_0(C_1) \cap b_2(C_2) = \emptyset$; (c) $b_1(C_1) \cap b_1(C_2) = \emptyset$; (d) $b_2(C_1) \cap b_2(C_2) = \emptyset$.

Proof. In the following discussion, $R_{i\theta}$ denotes the set of placements in C_i at orientation θ . (a) If $(X, \theta) \in b_0(C_1) \cap b_1(C_2)$ then $X \in \text{interior}(\Phi_1(C_2))$ and $X \in b(R_{1\theta})$. But this means that a small neighborhood of X will yield a point in common between $\text{interior}(\Phi_1(C_2))$ and $R_{1\theta}$. This contradicts the fact that $\Phi_1(C_2)$ should witness the criticality of θ , so $R_{1\theta}$ should do likewise (if $\Phi_1(C_2)$ intersects two maximal noncritical regions, then both of them witness the criticality of θ). (b) Essentially the same. (c) Let (X, θ) be in the intersection. Then $X \in \text{interior}(\Phi_1(C_1)) \cap \text{interior}(\Phi_1(C_2))$. Since $R_{1\theta}$ and $R_{2\theta}$ are continuous in θ , there is a θ' slightly greater than θ such that $X \in \text{interior}(R_{1\theta'}) \cap \text{interior}(R_{2\theta'})$ and θ' is in the range of both cells. But this implies $(X, \theta') \in C_1 \cap C_2$, so $C_1 = C_2$. (d) Essentially the same as (c). \square

Corollary 2.16. Only the following combinations of boundary parts yield adjacencies between two distinct cells C_1 and C_2 :

- (a) $b_0(C_1) \cap b_0(C_2) \neq \emptyset$, i.e., there exists θ in the common range of C_1 and C_2 such that $b(R_{1\theta}) \cap b(R_{2\theta}) \neq \emptyset$;
- (b) $b_1(C_1) \cap b_2(C_2) \neq \emptyset$, i.e., $\text{interior}(\Phi_1(C_1)) \cap \text{interior}(\Phi_2(C_2)) \neq \emptyset$ and the lower θ -bound of C_1 coincides with the upper θ -bound of C_2 ; or (vice-versa)
- (c) $b_1(C_2) \cap b_2(C_1) \neq \emptyset$, i.e., $\text{interior}(\Phi_1(C_2)) \cap \text{interior}(\Phi_2(C_1)) \neq \emptyset$.

Lemma 2.17. The relation $b(R_{1\theta}) \cap b(R_{2\theta}) \neq \emptyset$ remains invariant for all θ in the intersection of the ranges of C_1 and C_2 .

Proof. (Sketch). If θ is noncritical for both cells $C(R_{i\theta})$, and (X, θ) is a point common to their free boundary, then it is easy to show that X is in a unique segment $J(\theta)$ of the line at orientation θ passing through a fixed (real or virtual) corner v (see Figure 3), except for the situation depicted in Figure 3d. Thus the boundary intersection is nonempty whenever $J(\theta)$ is defined and has length greater than d . This implies, by straightforward continuity arguments, that the set of orientations θ , for which the given intersection is nonempty, is open. Let Θ represent this set of orientations. Similarly, if θ' is a limit orientation of such orientations θ , and θ' is noncritical for either cell, then $J(\theta')$ is defined and has length greater than d , so the given intersection is nonempty. Thus the set Θ is closed (as a subset of the common range of the two cells). Thus it is both open and closed within the common range. Being a nonempty subset of a connected range of orientations, it therefore coincides with that range. The case of Figure 3d can be handled separately. \square

Corollary 2.18. CG_{θ} is invariant between critical orientations. To be more precise, if θ_1 and θ_2 are such that every orientation in $[\theta_1, \theta_2]$ is noncritical, CG_{θ_1} is naturally isomorphic to CG_{θ_2} . \square

Definition. The *three-dimensional connectivity graph* CG is the graph whose vertices are the cells of FP and in which two cells C_1 and C_2 are adjacent if and only if either

$$(a) \ b_0(C_1) \cap b_0(C_2) \neq \emptyset,$$

$$(b) \ b_1(C_1) \cap b_2(C_2) \neq \emptyset, \text{ or}$$

$$(c) \ b_1(C_2) \cap b_2(C_1) \neq \emptyset.$$

Clearly these three cases are mutually exclusive.

In the first case we say that a *translational* adjacency exists between C_1 and C_2 , in the other two cases the adjacency is said to be *rotational*. The terms are derived from the fact that a *translational* adjacency allows the forearm to be translated from one cell to the other keeping the orientation (which must of course be in the range of both cells) fixed, as in the two-dimensional case. A *rotational* adjacency allows the forearm to cross the boundary by a small rotation without moving P . The definitions are illustrated in Figures 7 and 8. The preceding discussion allows us to state the following result. No proof is given, since it can be proved by arguments identical to those in Lemma 2.6.




Figure 7.




Figure 8.

Lemma 2.19. Z_1 is reachable from Z_2 by a continuous motion in FP iff $C(Z_1)$ is reachable from $C(Z_2)$ in CG . \square

Note. By Lemma 2.10, if $\Phi_1(C)$ has nonempty interior R , then either R is a maximal noncritical region or the union of two such regions with the free boundary separating them. In particular, a cell can be adjacent to zero, one, or two other cells at its lower (upper) θ -bound.

The above results allow us to describe fully the adjacencies of a cell C_1 in CG as follows:

- (1) If $\text{interior}(\Phi_1(C_1)) \neq \emptyset$, two cases are possible. Let θ be the critical orientation defining the lower θ -bound of C_1 . Let l be the line with orientation θ which manifests the criticality of θ , so $l \cap \Phi_1(C_1) \neq \emptyset$. If $\text{interior}(\Phi_1(C_1)) \cap l = \emptyset$, $\text{interior}(\Phi_1(C_1))$ is a maximal noncritical region with respect to θ (by Lemma 2.10 assuming the obstacles are in general position), so C_1 is rotationally adjacent to the unique cell C_2 for which $b_1(C_1) \subseteq b_2(C_2)$, i.e., for which $\text{interior}(\Phi_1(C_1)) \subseteq \text{interior}(\Phi_2(C_2))$ and whose upper θ -bound is θ (by Note following Lemma 2.19). On the other hand, if l cuts $\text{interior}(\Phi_1(C_1))$ into two regions, then C_1 is rotationally adjacent (at θ) to $C(R_1)$ and $C(R_2)$, where R_i are those two maximal noncritical regions which intersect the interior of $\Phi_1(C_1)$. (See Note following Lemma 2.10, which justifies the existence of these two cells.) A similar analysis holds for $\Phi_2(C_1)$.
- (2) Consider a translational adjacency of C_1 and C_2 at orientation θ ; thus $b_0(C_1) \cap b_0(C_2) \neq \emptyset$. Moreover, $b(R_{1\theta}) \cap b(R_{2\theta}) \neq \emptyset$, so $R_{1\theta}$ is adjacent to $R_{2\theta}$ in CG_θ , and by Corollary 2.18 $R_{1\theta'}$ will be adjacent to $R_{2\theta'}$ for all θ' in the common (angular) range of C_1 and C_2 . Thus all translational adjacencies of C_1 at orientation θ are valid in the range of orientations common to C_1 and all its "translational neighbors" at θ . As soon as one of the neighbors, say C_2 , disappears by virtue of θ leaving θ -range of C_2 , the corresponding translational adjacency disappears, possibly to

be replaced by an adjacency of C_1 to a new cell C_3 whose underlying region $R_{3\theta}$ is adjacent to that of C_1 at the current orientation.

The above description shows that, given CG_θ for every noncritical orientation, it is possible to build C_1 by

- (1) “opening” it at its lower θ -bound θ_1 and, if $\Phi_1(C_1)$ has non-empty interior, creating one or two rotational links.
- (2) adding initial $O(1)$ translational adjacencies from $CG_{\theta_1-\epsilon}$, for sufficiently small ϵ ,
- (3) updating its translational adjacencies according to $CG_{\theta-\epsilon}$ after each criticality θ brought about by creation or destruction of one of its translational neighbors,
- (4) and finally “closing” C at its upper θ -bound θ_1 , creating one or two rotational links if $\Phi_2(C_1)$ has non-empty interior.

The preceding paragraph is an informal description of the progress of the algorithm for building CG when restricted to its effect on a single cell. Note that knowledge of CG_θ for every noncritical orientation is equivalent by Corollary 2.18 to knowledge of CG_θ for each noncritical interval between two critical orientations.

3. Algorithmic details.

The connectivity graph is constructed following the model of the Leven-Sharir ‘ladder’ algorithm. Our construction differs in important details from theirs because the robot forearm interacts differently with real and virtual walls. The construction proceeds to first determine all the critical orientations, then compute the two-dimensional connectivity graph for a fixed orientation, and finally build the three-dimensional connectivity graph incrementally from the two-dimensional one by making the necessary adjustments in connectivity patterns as the orientation crosses each criticality.

3.1. Locating all critical orientations.

During this stage we process each corner v (whether real or virtual) and determine all critical orientations whose witnesses have v in their label. The following lemma shows that this is sufficient to identify all the criticalities.

Lemma 3.1. If θ is a critical orientation, then the ambiguous label of any witness region R for θ contains at least one corner.

Proof. Suppose that the left stop of R is ambiguous. Let M_θ (cf. Figure 9) be the maximal closed segment (at orientation θ) not penetrating obstacles and containing the left side of R (which may be degenerate); M_θ may touch obstacles and contain walls. Since θ is a critical orientation and R is a witness region, the following condition holds twice (independently): M_θ either passes through a corner or cuts two walls at points exactly d apart. Therefore, M_θ




Figure 9. Generic M_θ .

has to (i) pass through two corners, or (ii) pass through one such point and cut two walls at points d apart, or (iii) cut through two pairs of walls generating two pairs of points each pair exactly d apart.

However, (iii) may be excluded by assumptions of general position, and in the other two cases M_θ contains a (real or virtual) corner; this concludes the proof. \square

Fix a corner v . Note that all criticalities relevant to it can be divided into three classes:

- (1) those defined by another corner that is either visible from v or on a common wall with v ,
- (2) those that are defined by a line through v whose intersections with two walls are exactly d units apart, and v does not meet either wall, and
- (3) those defined by a line through v which meets another real wall at distance d from v .

The following computations will make use of the maximal star-shaped subset A of the free room V centered at v and a similar subset B of the workspace F . A can be easily computed from V by a sweeping technique in time $O(n \log n)$, while B can be computed in by a linear scan of (the boundary of) F . Both A and B will be represented by circular lists of their boundary edges.

Let us consider criticalities of each class separately.

- (1) It is easy to compute the criticalities of the first type. By traversing A , find the list of all corners visible from v (those are the vertices of A that are corners) including the other endpoints of walls incident to v . By the assumption of general position, no three corners are on the same line. Using brute force, compute for each corner w on the list the line-segment M_θ (where θ is the orientation of vw : it is potentially a critical orientation) and the set of virtual walls intersecting M_θ . The amount of data is clearly $O(n)$ per potential criticality θ , and this type of criticality can in fact have complexity $O(n)$, as illustrated in Figure 10. Let us see how to verify that θ is indeed a criticality, i.e., that




Figure 10.

M_θ contains a side of a witness region R for θ . First separate M_θ into its intervals of intersection with the workspace F by intersecting it with all the virtual walls it crosses (there may be none). Note that this does not require sorting as virtual walls are naturally ordered around O . For each segment I of the intersection $M_\theta \cap F$, construct segment J by extending I upwards by a segment of length d and clipping if necessary to ensure J does not protrude upwards beyond M_θ . In order for I to be the side of a witness region R the following criteria are necessary and sufficient: (i) $|J| > d$ (the case $|J| = d$ will be taken care of in (2) and (3) below) and (ii) both v and w be in the closure of J . These criteria can be verified in constant time per segment I , and therefore we can obtain all witness regions to θ being a criticality of the type indicated in overall time $O(n)$ per potential criticality θ . Thus we can compute $O(n)$ criticalities (per corner v) of complexity $O(n)$ in time $O(n^2)$, yielding a total of $O(n^2)$ $O(n)$ -criticalities computed in $O(n^3)$ time.

- (2) We will now turn our attention to the criticalities generated by a line passing through a corner v and cutting two walls a distance d apart. Consider three different cases:

- (i) The two walls in question are real. In this case we follow the procedure of [LS] by merging A with itself shifted by π and then examining the pair of real walls corresponding to each subinterval produced to find for which orientations the two are exactly d apart. At such orientations the line-segment M_θ has length d . To verify that it witnesses a critical orientation one must check only that its lower end is in the closure of the workspace F .
- (ii) The lower of the two walls is virtual and the upper is real: no other combination involving a virtual wall leads to a criticality. First we derive those criticalities where the real and virtual walls are on the same side of v . Consider the cyclically ordered list A of visible real-wall segments around v and compare it to the boundary of the maximal star-shaped subset B of F centered at v . By a merging operation one can determine all orientations θ such that the ray at orientation θ from v meets such a virtual wall w_1 on B and at distance d beyond it a real wall w_2 (in A). These orientations are the critical orientations of the kind we expect. Note that there are $O(n^2)$ such orientations for each v and each has complexity $O(1)$.
- (iii) Next we cover the subcase where the real and virtual walls are on opposite sides of v . Consider a triangular sector of A defined by v and an edge of A . Using brute force, in linear time compute the intersections of this sector with *all* virtual walls. Combining this information with the list of edges of A shifted by π , determine all orientations θ such that the line-segment $vX \subset V$ in direction θ connecting v to a *real* wall has length $a < d$ and the line in direction $\theta + \pi$ meets a virtual wall at a point Y such that the line segments vY has length $d - a$. This can be done by simultaneously sweeping through the lists described above, yielding $O(n^2)$ pairs of such points X and Y and thus computing another class of $O(n^2)$ criticalities of complexity $O(1)$ in time $O(n^2)$. This bound is tight as illustrated by Figure 11.




Figure 11.

- (3) Similar to (2iii). Identify each direction θ where a nearest real wall is met at a distance d from v ; a brute-force search through the virtual walls that are cut by a ray from v with orientation $\theta + \pi$ gives the desired witness region (the real wall at which the said ray leaves V is one of the candidates for the search).

3.2. Construction of CG_θ for a noncritical orientation θ .

Let θ be a particular noncritical orientation (again visualized as directed vertically upwards). Given the boundaries of the surrounding room and enclosed obstacles, it is required to partition the interior of the room outside the obstacles into maximal noncritical regions with respect to θ and to compute the adjacency relation on the regions explicitly. Let

VP denote the free room V partitioned by vertical line-segments (i.e., line-segments at orientation θ) through the corners of V : a vertical segment through a corner extends up to the walls directly above and below it. (Since θ is assumed noncritical, these line-segments can be characterized more precisely as follows: for each corner v its corresponding line-segment is the maximal vertical interval containing v in the closure of V .) It is easy to construct VP naively in time $O(n^2)$. The regions in VP are triangular or trapezoidal in shape. Each such region is then pruned by taking a line parallel to its upper segment, at distance d below it, and discarding that part of the region on or above this line (this may delete the whole region). Each remaining nonempty region R is a maximal noncritical region in the sense of [LS]; any vertical placement of the ladder with P in R is a free placement and the vertical ends of R represent either nonfree placements or a change in upper/lower stop. We can regard these regions as nodes in a graph G whose edges connect adjacent regions with non-trivial common (vertical) boundary.

Next compute the set VW of virtual walls. This can be done naively in $O(n^2)$ time, say, by sorting the corners of V in cyclic order around O and, rotating a sweepline around O , computing the subsegments of real walls visible from O . The virtual walls then are whatever radial edges are needed to connect these real wall-segments into the visibility polygon around O . This yields $O(n)$ virtual walls. Note that since each open triangular or trapezoidal region T of VP contains no wall, it contains no endpoint of any virtual wall. Assuming that VW is sorted cyclically around O , the set of virtual walls intersecting the interior of such a trapezoidal region T can be identified in $O(n)$ time, in cyclic order around O : these intersecting walls partition T radially into several strips, alternately inside and outside the workspace F .

Let U be such a strip. Being the intersection of T with a wedge whose apex is at O , U is bounded by at most two virtual walls, and therefore has between four and six sides. Also, U has at most two corners where a virtual wall meets one of the non-vertical sides of T : in other words, U possesses at most two corners which are virtual corners of F . Vertical lines

through these corners partition U into at most three open trapezoidal regions, and by definition, the upper and lower segments bounding each such regions contain no real or virtual corner. Therefore each such region is entirely contained in some maximal noncritical region of CG_θ . There are $O(n^2)$ such regions. We can regard these regions as being nodes in a graph G' whose edges connect adjacent regions possessing a nontrivial common vertical boundary. Then G' can be regarded as a size $O(n^2)$ subdivision of CG_θ , and it can have been constructed in $O(n^2)$ time overall. This bound is tight as illustrated in Figure 5.

Finally the graph G' may contain chains of contiguous regions all possessing the same upper and lower stops (this can arise where vertices separating maximal noncritical regions in the sense of [LS] are occluded by some vertical wall). It is easy to merge such chains of vertices together, and in this way the desired connectivity graph CG_θ is constructed. The procedure for computing CG_θ is illustrated in Figure 12.

3.3. Incremental construction of CG from CG_θ .

This section describes the final step of the algorithm — the construction of the connectivity graph CG which defines the partition of the space of all free placements FP into cells together with the connectivity pattern among them. The reachability problem is then reduced to identifying the cells where the initial and final placements are located and then determining whether the two are connected by at least one path in the graph CG . The algorithm is easily modified to create a path between starting and destination placements with no performance




Figure 12.

degradation, as a path connecting arbitrary points of two adjacent cells can be computed in constant time.

Again, our presentation follows [LS] with modifications necessary to take virtual walls into account. We will assume that a sorted list of all $O(n^3)$ critical orientations is available. It can be compiled by sorting in time $O(n^3 \log n)$. It is also assumed that each critical orientation carries with it the information on how it was created. This includes the *label* of the critical orientation, the corner that the line with such an orientation has to pass through to exhibit its criticality, the other corner (if any) through which it passes, the real walls at which it terminates, and the virtual walls bounding witness regions which it meets.

The construction begins by choosing a noncritical orientation θ_0 and building the graph CG_{θ_0} in time $O(n^2)$ as described above. There are two structures maintained throughout the construction: CG_{θ} , the two-dimensional connectivity graph with respect to the current orientation θ , and an incomplete version of CG , which includes cells with orientation ranges truncated to within $[\theta_0, \theta]$ and adjacency relation restricted to those cells. Set the current orientation θ to be θ_0 . Initialize CG by copying CG_{θ_0} and modifying each node R (which is a maximal noncritical region with respect to θ_0) to represent instead a cell with R as its underlying region R_{θ} , θ_0 as lower θ -bound, and a special “unknown” value as the upper θ -bound. This will take $O(n^2)$ time, as a copy of CG_{θ_0} needs to be made.

As we have already noted the connectivity structure of CG_{θ} only changes when θ crosses a criticality (Corollary 2.18). When θ crosses a criticality only the witness regions of the criticality and their neighbors will be affected, and because CG_{θ} has bounded degree (assuming the obstacles are in general position, so, for example, each region in CG_{θ} for a noncritical θ shares its free left boundary with at most two other regions), this means that $O(k)$ work is needed to process the graph when crossing a criticality of complexity k . The changes needed can easily be determined by examining the description of the criticality alone (see Figure 13).




Figure 13.

Let us describe some of the details of how the graph CG is built in this way. At each critical orientation θ' , for each region R being destroyed in $CG_{\theta' - \epsilon}$ (i.e., for each R such that the upper θ -bound of $C(R)$ is θ' ; ϵ is chosen so there are no criticalities between $\theta' - \epsilon$ and θ'), make θ' the upper θ -bound of the corresponding "unfinished" cell in CG , and for each R being introduced into $CG_{\theta' - \epsilon}$ create a new cell in CG based on R with θ' as its lower θ -bound and unknown upper θ -bound with all translational adjacencies inherited from R . When a region in CG_{θ} is destroyed, it either contracts to a set with empty interior, or becomes one or two (adjacent) maximum noncritical regions with respect to θ' with ambiguous labels (Lemma 2.10). The latter type of Φ_2 creates a rotational adjacency per noncritical region with respect to θ' contained in Φ_2 (see Figure 8), for a total of at most two per cell per θ -bound. At this point all new rotational links associated with θ' are added. After processing is completed we have a 'split' connectivity graph whose range is $[\theta_0, \theta_0 + 2\pi]$: we can assume that the structure of CG_{θ_0} has been saved, and the processing has yielded (as a by-product) a connectivity graph $CG_{\theta_0 + 2\pi}$ which is of course naturally isomorphic to the former graph. At this point we have to "repair" the cells split in two by θ_0 .

For simplicity let us assume that the lines joining all pairs of corners in the workspace F have slopes different from θ_0 : these slopes were generated when locating criticalities, so this causes no problems (some but not necessarily all of these slopes are indeed critical orientations). Let G abbreviate the initial translation graph CG_{θ_0} and G' the final translation graph $CG_{\theta_0 + 2\pi}$. As remarked, G and G' match exactly. Each node in these graphs may be

assumed labeled with its four stops. Again visualizing the direction θ_0 as vertical, let V and V' be the node-sets in each graph — suitably represented, and including the coordinates position of their lower left corners; sort the combined set according to the y -coordinates within x -coordinates of their lower left corners. Then corresponding nodes in each graph will be adjacent in the resulting sorted list. Thus the nodes can be matched up, and the corresponding half-cells merged. The overall work for this phase is $O(n^2 \log(n))$.

Thus CG has been computed. To solve a given motion-planning query, specified by initial and final placements Z_0 and Z_1 , it is enough to compute the labels of $C(Z_0)$ and $C(Z_1)$ and locate the corresponding cells in CG , perhaps by a linear scan of its cells. The problem can then be solved by a straightforward search of CG . This concludes our description of the final stage of the $O(n^3 \log n)$ algorithm that reduces path planning for a specific model of a two-dimensional robot to searching of a graph of size $O(n^3)$.

4. Concluding remarks.

We have presented an exact algorithm to solve motion-planning problems for an idealized two-link planar robot arm. The main effort has been to analyze the configuration space of such an arm: in contrast with the related 'ladder mover's' problem studied in the literature, the configuration space complexity is $O(n^3)$. Our algorithm computes it with a small extra overhead of $\log(n)$. The analysis leaves some interesting questions open:

How realistic is the model? It seems to us that the most unrealistic assumption is that of the infinite extensibility of the 'upper arm.' Minimal and maximal radii should probably be prescribed for the length of this link. It does not seem too difficult to incorporate this restriction into the algorithm: the effect is to confine the workspace in an annular planar region, so the only variant, essentially, is accommodating circular arcs in the description of the workspace. Similarly, limiting the range of angles through which the upper arm can rotate about the shoulder obviously causes no difficulty: everything is confined between two extra virtual walls. To limit the range of rotation at the elbow is easy if the range of rotation

is absolute, i.e., related to the plane rather than the mutual angle between the two links at the elbow, since that would only involve setting limits on the angular range of the graph CG . To incorporate *relative* angular restraints at the elbow is a different matter and it is not clear that the method of this paper is appropriate for such a situation. Another variation of the problem would be to model the upper arm not as a telescoping link but a fixed-length sliding boom (cf. Fortune, Wilfong and Yap [FWY]). This variation may be amenable to the methods discussed here.

It is not clear whether our method generates 'efficient' solutions to the motion-planning problem, i.e., solution paths formed from sufficiently few 'primitive path-segments' (such as translations between adjacent cells). It is possible, but seems unlikely, that nonlinear lower bounds such as Ke and O'Rourke's [KO] can be shown to hold here: indeed, we conjecture that while the number of components of configuration space can be $\Omega(n^3)$, any two mutually accessible placements can be connected by a path of descriptive complexity $O(n)$. It would seem that the constraint imposed by the upper arm can both subdivide (which we know) configuration space but yet produce simpler components: for instance, every such component is of course simply connected.

5. References.

- [FWY] S. Fortune, G. Wilfong, and C.-K. Yap, "Coordinated Motion of Two Robot Arms", *Proc. IEEE International Conference on Robotics and Automation*, pp. 1216-1223, August 1986, San Francisco.
- [KO] Y. Ke and J. O'Rourke, "Lower Bounds on Moving a Ladder in Two and Three Dimensions", *Proceedings of the 3rd Annual ACM Symposium on Computational Geometry*, pp. 136-146, June 1987, Waterloo.
- [LS] D. Leven and M. Sharir, "An Efficient and Simple Motion Planning Algorithm for a Ladder Amidst Polygonal Barriers", *Journal of Algorithms*, Vol. 8, 1987, pp. 192-215.

- [SS1] J.T. Schwartz and M. Sharir, "On the Piano Movers' Problem: I. The Case of a Two Dimensional Rigid Body Moving Amidst Polygonal Barriers", *Comm. Pure Appl. Math.*, Vol. 36, 1983, pp. 345-398.
- [SS2] J.T. Schwartz and M. Sharir, "On the Piano Movers' Problem: II. General Techniques for Computing Topological Properties of Real Algebraic Manifolds", *Advances in Applied Mathematics*, Vol. 4, 1983, pp. 298-351.
- [PS] F.P. Preparata and M.I. Shamos, *Computational Geometry: An Introduction*, Springer-Verlag, New York, 1985.

NYU COMPSCI TR-314 c.1
Aronov, Boris
Analysis of the motion-
planning problem for a
simple two-link planar arm

- NYU COMPSCI TR-314 c.1 --
- Aronov, Boris --
- Analysis of the motion-
- planning problem for a --
= simple two-link planar arm ==

This book may be kept

FOURTEEN DAYS

A fine will be charged for each day the book is kept overtime.

